

Smart machines and the SP theory of intelligence

J Gerard Wolff*

February 22, 2014

Abstract

These notes describe how the *SP theory of intelligence*, and its embodiment in the *SP machine*, may help to realise cognitive computing, as described in the book *Smart Machines*. In the SP system, information compression and a concept of *multiple alignment* are centre stage. The system is designed to integrate such things as unsupervised learning, pattern recognition, probabilistic reasoning, and more. It may help to overcome the problem of variety in big data, it may serve in pattern recognition and in the unsupervised learning of structure in data, and it may facilitate the management and transmission of big data. There is potential, via information compression, for substantial gains in computational efficiency, especially in the use of energy. The SP system may help to realise data-centric computing, perhaps via a development of Hebb's concept of a 'cell assembly', or via the use of light or DNA for the processing of information. It has potential in the management of errors and uncertainty in data, in medical diagnosis, in processing streams of data, and in promoting adaptability in robots.

Keywords: unsupervised learning, big data, computational efficiency, data-centric computing, veracity, uncertainty.

1 Introduction

The book *Smart Machines: IBM's Watson and the Era of Cognitive Computing* [3] provides a very interesting account of the need for "a new generation

*Dr Gerry Wolff, BA (Cantab), PhD (Wales), CEng, MBCS (CITP); CognitionResearch.org, Menai Bridge, UK; jgw@cognitionresearch.org; +44 (0) 1248 712962; +44 (0) 7746 290775; *Skype*: gerry.wolff; *Web*: www.cognitionresearch.org.

of tools—cognitive technologies—that help us to penetrate complexity and comprehend the world around us so that we can make better decisions and live more successfully and sustainably.” (*ibid.*, Preface).

These notes are about how the *SP theory of intelligence* and its embodiment in the *SP machine* may help to translate that vision into reality.

I’ll first introduce the SP concepts briefly and then try to show how they may facilitate some of the possibilities described in *Smart Machines*.

2 Introduction to the SP theory and SP machine

The SP theory, which has been under development for several years, aims to simplify and integrate concepts across artificial intelligence, mainstream computing and human perception and cognition, with information compression as a unifying theme.

The theory is conceived as an abstract brain-like system that, in an ‘input’ perspective, may receive ‘New’ information via its senses, and compress some or all of it to create ‘Old’ information, as illustrated schematically in Figure 1. In the theory, information compression is the mechanism both for the learning and organisation of knowledge and for pattern recognition, reasoning, problem-solving, and more.

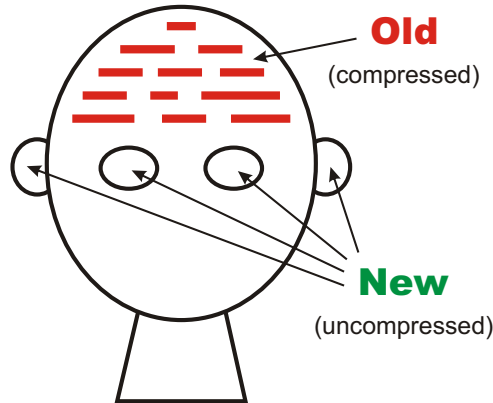


Figure 1: Schematic representation of the SP system from an ‘input’ perspective.

In the SP theory, all kinds of knowledge are represented with *patterns*: arrays of atomic symbols in one or two dimensions. A key part of the system

is a concept of *multiple alignment*, like that shown in Figure 2.¹

0				t	h	e					a	p	p	l	e	s				a	r	e			s	w	e	e	t	0
1																														1
2																														2
3																														3
4																														4
5																														5
6	S	Num																												6
7																														7
8	Num	PL	;																											8

Figure 2: A multiple alignment created by the SP computer model that achieves the effect of parsing a sentence (“t h e a p p l e s a r e s w e e t”).

The SP theory is realised in a computer model, SP70, which may be regarded as a first version of the SP machine. It is envisaged that the SP computer model will provide the basis for the development of a high-parallel, open-source version of the SP machine, as shown schematically in Figure 3. This will be a means for researchers everywhere to explore what can be done with the system and to create new versions of it.

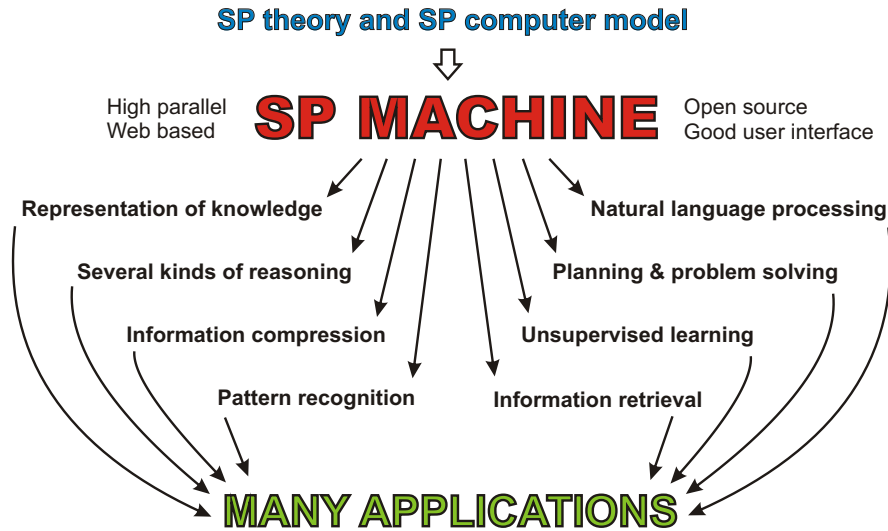


Figure 3: Schematic representation of the development and application of the SP machine.

¹The concept of multiple alignment in the SP system ([14, Section 4]; [10, Section 3.4]) is borrowed from that concept in bioinformatics, but with important differences.

The SP system provides a model for several aspects of computing and cognition, including unsupervised learning, concepts of computing, aspects of mathematics and logic, the representation of knowledge, natural language processing, pattern recognition, several kinds of reasoning, information storage and retrieval, planning and problem solving, and aspects of neuroscience and of human perception and cognition [14, 10].

There is a relatively full account of the SP system in [10], an extended overview in [14], an account of its existing and expected benefits and applications in [15], a description of its foundations in [12], and an introduction to the system in [13]. There is more information in www.cognitionresearch.org/sp.htm.

3 Unsupervised learning

“While traditional computers must be programmed by humans to perform specific tasks, cognitive systems will learn from their interactions with data and humans and be able to, in a sense, program themselves to perform new tasks.” [3, p. 7].

The SP programme derives in part from earlier research developing computer models of language learning.²

At a fundamental level—compression of information via the matching and unification of patterns—the SP computer model works in the same way as the earlier models of language learning. Like the earlier models, the SP computer model is able to discover generative grammars from data, including segmental structure, classes of structure, and abstract patterns.

But in developing the SP system, a radical reorganisation has been needed to meet the goal of simplifying and integrating concepts across artificial intelligence, mainstream computing, and human perception and cognition. Unlike the earlier models, multiple alignment is central in the workings of the SP computer model, including unsupervised learning. A bonus of the new structure is potential for the unsupervised learning of class hierarchies, part-whole hierarchies, and discontinuous dependencies in data.

The emphasis on unsupervised learning in the SP programme of research is because of evidence that human learning, including the learning of natural languages, does not depend on the provision of a ‘teacher’, or anything equivalent—although such assistance may be helpful.

A strength of the SP system is that it can discover significant entities, classes of entity, and other structures in data, not merely statistical asso-

²See bit.ly/JCd6jm.

ciations between pre-established structures. An important idea in the SP programme is the ‘DONSVIC’ principle [14, Section 5.2]: the conjecture, supported by evidence, that information compression is the key to the discovery of ‘natural’ structures, meaning the kinds of things that people naturally recognise, such as words, objects, and classes of objects.

Learning in the SP system has much of the flavour of:

“[building and maintaining] a giant jigsaw puzzle on the fly, making connections between newly acquired data and older information.” [3, p. 55].

and of:

“... finding many needles in a field full of haystacks, combining those needles, then searching the smaller but still substantial collection for the needles made of gold—the truly valuable slivers of insight.” [3, p. 56].

4 Big data

“The emergence of social networking, sensor networks, and huge storehouses of business, scientific, and government records creates an abundance of information tech-industry insiders call ‘big data’. ... This digital universe is growing at about 60 percent each year.” [3, p. 5].

“We need a new generation of data storage, management, and analytics tools that will improve our ability to gather, meld, and make sense of huge amounts of data and, in some cases, to perform a complex ballet of tasks in real time.” [3, p. 45].

The SP system may assist in the management and analysis of big data in four main ways:

- *Helping to overcome the problem of variety in big data.*

“The manipulation and integration of heterogeneous data from different sources into a meaningful common representation is a major challenge.” [6, p. 76].³

³See also the problem of variety described in [4] and [3, p. 45]

The versatility of SP patterns, within the multiple alignment framework, in the representation and processing of diverse kinds of knowledge ([14, Sections 5 to 13]; [10, Chapters 5 to 10]) makes them a contender for the role of “meaningful common representation” or universal format for knowledge. There is potential for the translation of heterogeneous data into that universal format, for harmonisation in the representation of different kinds of knowledge, and for automatic structuring of knowledge via processes for unsupervised learning (Section 3).

- *Interpretation.* In the interpretation of data, the SP system has capabilities that include such things as natural language processing, several kinds of reasoning, planning, and problem solving. It has several strengths in pattern recognition, where recognition may be achieved: at multiple levels of abstraction; with “family resemblance” or polythetic categories; in the face of errors of omission, commission or substitution in data; with the calculation of a probability for any given identification, classification or associated inference; with sensitivity to context in recognition; and with the seamless integration of pattern recognition with other aspects of intelligence—reasoning, learning, problem solving, and so on ([14, Section 9]; [10, Chapter 6]).
- *Unsupervised learning.* As outlined in Section 3, unsupervised learning is a key feature of the SP system. With big data, there is clear potential for the previously-mentioned DONSVIC principle: *the discovery of natural structures via information compression* [14, Section 5.2].
- *Information compression.* Information compression is central in how the SP system works. Potential benefits with big data include:
 - Reducing the volume of data and thus facilitating its storage and management. In that connection, there are reasons to believe that the system may achieve higher levels of compression than conventional systems for information compression [15, Section 6.7].
 - By sending only ‘code’ and not ‘grammar’, there is potential for very substantial economies in the transmission of big data [15, Section 6.7.1]. There is more about this in Section 5.3.
 - It appears that information compression can mean substantial gains in computational efficiency, especially in the use of energy (Section 5).
 - In the SP system, information compression provides a key to the management of errors and uncertainty in data (Section 7).

- Knowledge structures created by the system, including class-inclusion hierarchies, part-whole hierarchies, and their integration [14, Section 9.1], may facilitate retrieval of information by serving, in effect, as hierarchical indices.

How the SP system may assist in the management and analysis of big data will be described more fully in a new article [7].

5 Computational efficiency via information compression

“... we’re reaching the limits of our ability to make [gains in the capabilities of CPUs] at a time when we need even more computing power to deal with complexity and big data. And that’s putting unbearable demands on today’s computing technologies—mainly because today’s computers require so much energy to perform their work.” [3, p. 9].

“The human brain is a marvel. A mere 20 watts of energy are required to power the 22 billion neurons in a brain that’s roughly the size of a grapefruit. To field a conventional computer with comparable cognitive capacity would require gigawatts of electricity and a machine the size of a football field. So, clearly, something has to change fundamentally in computing for sensing machines to help us make use of the millions of hours of video, billions of photographs, and countless sensory signals that surround us. ... Unless we can make computers many orders of magnitude more energy efficient, we’re not going to be able to use them extensively as our intelligent assistants.” [3, p. 75, p. 88].

In the quest for greater efficiency in processing, especially energy efficiency, the SP system may make what is potentially a very substantial contribution via information compression (this section) and via data-centric computing (Section 6).

Since information processing in the SP system means compression of information via the matching and unification of patterns, *anything that increases the efficiency of searching for good full and partial matches between patterns will also increase the efficiency of information processing.*

It appears that information compression can itself be a means of increasing the efficiency of searching, as described in the next two subsections.

Information compression may also yield economies in the transmission of information (Section 5.3, below).

5.1 Reducing the sizes of data to be searched and of search terms

As described in [15, Section 6.7.2], if we wish to search a body of information, I , for instances of a pattern like “Treaty on the Functioning of the European Union” the efficiency of searching may be increased:

- By reducing the size of I so that there is less to be searched. The size of I may be reduced by replacing all but one of the instances of “Treaty on the Functioning of the European Union” with a relatively short code like “TFEU”, and likewise other recurrent patterns. More generally, the size of I may be reduced via the compression processes in the SP system.
- By searching with a short code like “TFEU” instead of a relatively large pattern like “Treaty on the Functioning of the European Union”. Other things being equal, a smaller search pattern means more efficient searching.

With regard to the second point, there is potential to cut out some searching altogether by “hard wiring” the connection between each instance of a code (“TFEU” in this example) and the thing that it represents (“Treaty on the Functioning of the European Union”). In *SP-neural* (Section 6.1), there are connections of that kind between “pattern assemblies”, as shown schematically in Figure 4.

5.2 Concentrating search where good results are most likely to be found

If we want to find some strawberry jam, our search is more likely to be successful in a supermarket than it would be in an antiques shop or a showroom for second-hand cars. This may seem too simple and obvious to deserve comment but it illustrates the extraordinary knowledge that most people have of an informal ‘statistics’ of the world that we inhabit, and how that knowledge may help us to minimise effort.⁴

⁴See also G. K. Zipf’s *Human Behaviour and the Principle of Least Effort* [16].

Where does that statistical knowledge come from? In the SP theory, it flows directly from the central role of information compression in our perceptions, learning and thinking, and from the intimate relationship between information compression and concepts of prediction and probability [5].

Although the SP computer model calculates probabilities for some purposes (see Section 7), it actually uses levels of information compression as a guide to search. Those levels are used, with heuristic search methods (including escape from ‘local peaks’), to ensure that searching is concentrated in areas where it is most likely to be fruitful [10, Sections 3.9, 3.10, and 9.2]. This not only speeds up processing but yields Big-O values for computational complexity that are within acceptable limits [10, Sections 3.10.6, 9.3.1, and A.4].

5.3 Economies in the transmission of data

“To control costs, designers of the [DOME] computing system have to figure out how to minimize the amount of energy used for processing data. At the same time, since so much of the energy in computing is required to move data around, they have to discover ways to move the data as little as possible.” [3, p. 65].

Although this quote may refer in part to movements of data such as those between the CPU and the memory of a computer, the discussion here is about transmission of data over longer distances such as, for example, via the internet.

As was mentioned in Section 4, there is potential with the SP system for very substantial economies in the transmission of data [15, Section 6.7.1]. Any body of data, I , may be compressed by encoding it in terms a ‘grammar’ (G), provided that G contains the kinds of structures that are found in I (Section 7). Then I may be sent from A to B by sending only the ‘encoding’ (E). Provided that B has a copy of G , I may be recreated with complete fidelity by means of the SP system ([14, Section 4.5]; [10, Section 3.8]). Since E would normally be very much smaller than the I from which it was derived, it seems likely that there would be a net gain in efficiency, allowing for the computational costs of encoding and decoding.

Since a copy of G must be transmitted to B, any savings will be relatively small if it is used only for the decoding of a single instance of E . But significant savings are likely if, as would normally be the case, one copy of G may be used for the decoding of many different instances of E , representing many different bodies of information.

5.4 Potential gains in computational efficiency

No attempt has yet been made to quantify potential gains in computational efficiency from the compression of information, as described in Sections 5.1, 5.2, and 5.3, but they could be very substantial:

- Since information compression is fundamental in the workings of the SP system, there is potential for corresponding savings in all parts and levels in the system.
- The entire structure of knowledge that the system creates for itself is intrinsically statistical, with potential on many fronts for corresponding savings in computational costs and associated demands for energy.

It is anticipated that the proposed high-parallel version of the SP machine (Section 2) will provide a means of exploring these aspects of the system.

6 Data-centric computing

“What’s needed is a new architecture for computing, one that takes more inspiration from the human brain. Data processing should be distributed throughout the computing system rather than concentrated in a CPU. The processing and the memory should be closely integrated to reduce the shuttling of data and instructions back and forth.” [3, p. 9].

“Scientists at IBM Research believe that to make computing sustainable in the era of big data, we will need a different kind of machine—the data-centric computer. ... Machines will perform computations faster, make sense of large amounts of data, and be more energy efficient.” [3, p. 88].

The SP concepts may help to integrate processing and memory, as described in the next two subsections.

6.1 SP-neural

Although the main emphasis in the SP programme has been on developing an abstract framework for the representation and processing of knowledge, the theory includes proposals—called *SP-neural*—for how those abstract concepts may be realised with neurons [10, Chapter 11].

Figure 4 shows in outline how an SP-style conceptual structure would appear in SP-neural. It is envisaged that SP patterns would be realised with *pattern assemblies*—groupings of neurons like those shown in the figure within broken-line envelopes.

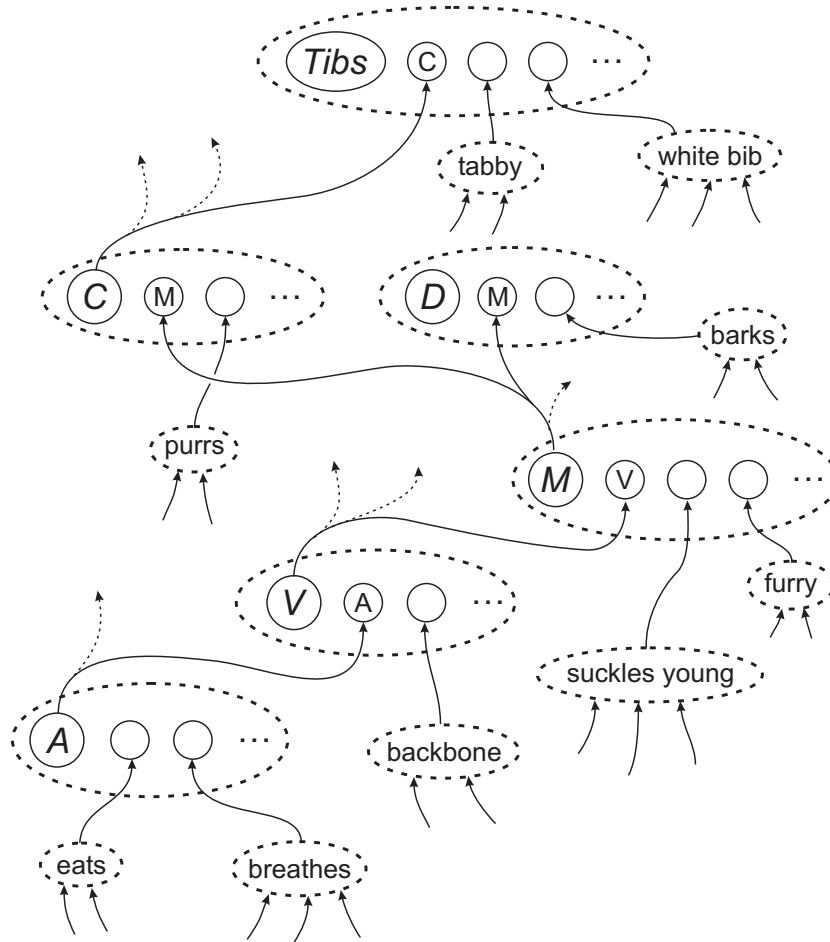


Figure 4: An example showing schematically how SP-neural may represent class-inclusion relations, part-whole relations, and their integration. *Key:* ‘C’ = cat, ‘D’ = dog, ‘M’ = mammal, ‘V’ = vertebrate, ‘A’ = animal, ‘...’ = further structure that would be shown in a more comprehensive example. Pattern assemblies are surrounded by broken lines and each neuron is represented by an unbroken circle or ellipse. Lines with arrows show connections between pattern assemblies and the flow of sensory signals in the process of recognising something (there may also be connections in the opposite direction to support the production of patterns). Connections between neurons within each pattern assembly are not marked.

The whole scheme is quite different from ‘artificial neural networks’ as they are commonly conceived in computer science.⁵ It may be seen as a development of Donald Hebb’s [2] concept of a ‘cell assembly’, with more precision about how structures may be shared, and other differences.⁶

In SP-neural, the structure of knowledge is reflected directly in groupings of neurons and their interconnections, as shown in the figure. Although the details have not been worked out, it is envisaged that such things as pattern recognition would be achieved via the transmission of impulses between pattern assemblies, and via the transmission of impulses between neurons within each pattern assembly. In keeping with what is known about the workings of brains and nervous systems, it is likely that there would be important roles for both excitatory and inhibitory signals.

In short, neurons in SP-neural serve for both the representation and processing of knowledge, with close integration of the two—in accordance with the concept of data-centric computing.

6.2 Computing with light or chemicals

The SP concepts appear to lend themselves to computing with light or chemicals, perhaps by-passing such things as transistors or logic gates that have been prominent in the development of electronic computers [15, Section 6.10.6].⁷

At the heart of the SP system is a process of finding good full and partial matches between patterns. This may be done with light, with the potential advantage that light beams may cross each other without interference. Another potential advantage is that, with collimated light, there may be relatively small losses over distance—although distances should probably be minimised to save on transmission times and to minimise the sizes of computing devices. There appears to be potential to create an optical or opti-

⁵See, for example, “Artificial neural network”, Wikipedia, en.wikipedia.org/wiki/Artificial_neural_network, retrieved 2013-12-23.

⁶In particular, unsupervised learning in the SP system ([14, Section 5]; [10, Chapter 9]) is radically different from the “Hebbian” concept of learning (see, for example, “Hebbian theory”, Wikipedia, http://en.wikipedia.org/wiki/Hebbian_learning, retrieved 2013-12-23), described by Hebb [2] and adopted as the mechanism for learning in most artificial neural networks. By contrast with Hebbian learning, the SP system, like a person, may learn from a single exposure to some situation or event. And, by contrast with Hebbian learning, it takes time to learn a language in the SP system because of the complexity of the search space, not because of any kind of gradual strengthening or “weighting” of links between neurons [10, Section 11.4.4].

⁷“The most promising means of moving data faster is by harnessing photonics, the generation, transmission, and processing of light waves.” [3, p. 93].

cal/electronic version of SP-neural.

Finding good full and partial matches between patterns may also, potentially, be done with chemicals such as DNA,⁸ with potential for high levels of parallelism, and with the attraction that DNA can be a means of storing information in a very compact form, and for very long periods [1].

With both optical and chemical versions of the SP system, there seems to be potential for achieving data-centric integration of knowledge and processing.

7 Errors and uncertainties in data

“Organizations face huge challenges as they attempt to get their arms around the complex interactions between natural and human-made systems. The enemy is uncertainty. In the past, since computing systems didn’t handle uncertainty well, the tendency was to pretend that it didn’t exist. Today, it is clear that that approach won’t work anymore. So rather than trying to eliminate uncertainty, people have to embrace it.” [3, pp. 50–51].

The SP system has potential in the management of errors and uncertainties in data. In summary:

- Owing to the previously-mentioned close connection between information compression and concepts of prediction and probability, the whole system is inherently probabilistic. Every SP pattern has an associated frequency of occurrence. From that information, a probability may be derived for each multiple alignment and for each of what are normally several inferences that may be derived from each multiple alignment ([14, Section 4.4]; [10, Section 3.7 and Chapter 7]).⁹
- Unsupervised learning via information compression provides a neat solution to the problem of “dirty data” in the learning of a natural language ([14, Section 5.3]; [8]): how it is that we can develop a keen sense of what does or does not belong in our native language or languages, despite the fact that most of the speech that children hear contains the kinds of haphazard errors that people make in talking, and in the face

⁸See, for example, “DNA computing”, Wikipedia, bit.ly/1gfEP4p, retrieved 2013-12-30.

⁹Although the SP system is fundamentally probabilistic, it can, if required, be constrained to yield all-or-nothing results, much as in a conventional computer [10, Chapter 10].

of evidence that language learning may be achieved without the benefit of error correction by a teacher, or anything equivalent.

In brief, the product of learning comprises a ‘grammar’ (G) and an ‘encoding’ (E) of the original data in terms of the grammar. The two together achieve lossless compression of the original data, including all the errors. But G , which may be seen to capture the essence of the target language, excludes everything which is rare—which is mostly those haphazard errors that people make in speaking (these are rare individually, although collectively they are quite common). Anything that is a little more frequently-occurring than rare may acquire the status of linguistic irregularity—such as ‘bought’ (not ‘buyed’) as the past tense of ‘buy’—or it may be seen as a dialect form.

These principles apply to any kind of data, not just linguistic data.

- In operations such as parsing natural language or pattern recognition, the SP system is robust in the face of errors of omission, of commission, or of substitution ([14, Section 4.2.2]; [10, Section 6.2.1]). In the same way that we can recognise things visually despite disturbances such as falling leaves or snow, or our own blinking (and likewise for other senses), the SP system can cope quite well with inputs that are not totally correct.

Although there are advantages in using probabilities, the flipside—for both people and machines—is that mistakes can be made. We may bet on “Dancing Queen” but find that “Kiss me Kate” is the winner. In the same way that people can be fooled by a frequently-repeated lie, any probabilistic machine will be vulnerable to systematic distortions in data.

8 Other possibilities

Outline here are some other aspects of the SP system, and how they may help to realise the possibilities described in *Smart Machines*.

8.1 Medical diagnosis

“One of the next challenges for Watson is to help doctors diagnose diseases and assess the best treatments for individual patients.”
[3, p. 2].

How the SP system may assist with medical diagnosis is described in [9].

8.2 Processing streams of data

“Increasingly, vitally important insights can be gained from analyzing information that’s on the move. ... Rather than placing the data in a database first, the computer analyzes it as it comes in from a variety of sources, continually refining its understanding of the data as conditions change. This is the way humans process information.” [3, pp. 49–50].

In its overall organisation, the SP system is designed to process streams of “New” information, as shown schematically in Figure 1, and very much in the spirit of the quotation above.

As indicated in Section 4, the SP system may also provide a handle on the problem of variety in data, mentioned in the quotation.

8.3 The need for radical innovations

“Soon incremental innovation will no longer be sufficient. ... We need more radical innovations.” [3, pp. 17–18].

The idea that all kinds of computing and cognition may be understood as information compression via multiple alignment is indeed a radical innovation. But it is not merely an eccentric blind alley. Detailed research has shown its potential in several areas with a bearing on the creation of smart machines.

8.4 Augmenting our senses

“... most images and audio on the web are searchable only via metadata—words that are manually typed into forms by humans. So its vital to develop systems that can recognize images and sounds more like humans do.” [3, p. 69].

As was mentioned in Section 4, the SP system has several strengths in pattern recognition (see also [15, Section 9]; [10, Chapter 6]). How the system may be applied to vision is described in [11]. Because of the generality and versatility of the system [15, Section 4], it can probably be applied in the recognition of sounds and other kinds of sense data.

8.5 Achieving human-like flexibility and adaptability in robots

“... as of now, robots remain firmly in the von Neumann computing paradigm. They must be programmed in advance by people to deal with nearly every situation they encounter.” [3, p. 72].

How, via the SP system, robots may achieve human-like flexibility and adaptability is outlined in [15, Section 6.3]. With robots, especially autonomous robots, the SP system has at least two other attractions: versatility in different aspects of intelligence; and potential for computational efficiency, especially energy efficiency (Sections 5 and 6).

8.6 The need for integration

“Today, as scientists labor to create machine technologies to augment our senses, there’s a strong tendency to view each sensory field in isolation as specialists focus only on a single sensory capability. Experts in each sense don’t read journals devoted to the others senses, and they don’t attend one another’s conferences. Even within IBM, our specialists in different sensing technologies don’t interact much. Yet if machines are to help humans understand the world, they have to make sense of it and communicate about it in a way that’s familiar and comprehensible to humans. This integration of data from various sensing technologies is beginning to happen in multimedia and visual analytics, where vision and sound are correlated. But that’s just the start of what will be required in the next era of computing.” [3, p. 74].

This quotation touches on an issue which is right at the heart of the SP programme of research: the need for simplification and integration of concepts across different fields. The SP theory aims to apply this principle across artificial intelligence, mainstream computing, and human perception and cognition. Broadening the context in this way improves the chances of creating a theory that, in accordance with Occam’s Razor, combines conceptual *simplicity* with descriptive or explanatory *power* [15, Section 2].

9 Conclusion

The SP system may help to realise cognitive computing, as described in *Smart Machines*. It is designed to simplify and integrate concepts across artificial

intelligence, mainstream computing, and human perception and cognition. It has potential for unsupervised learning of structure in data, and it may help in the management and analysis of big data. Information compression can mean substantial gains in computational efficiency, especially in the use of energy. The system may help to realise data-centric computing. And it may facilitate the management of errors and uncertainty in data.

The creation of a high-parallel, open-source version of the SP machine, as outlined in Section 2, would be a means for researchers everywhere to explore what can be done with the system and to create new versions of it.

References

- [1] N. Goldman, P. Bertone, S. Chen, C. Dessimoz, E. M. LeProust, B. Sipos, and E. Birney. Towards practical, high-capacity, low-maintenance information storage in synthesized DNA. *Nature*, 494(7435):77–80, 2013.
- [2] D. O. Hebb. *The Organization of Behaviour*. John Wiley & Sons, New York, 1949.
- [3] J. E. Kelly and S. Hamm. *Smart machines: IBM’s Watson and the era of cognitive computing*. Columbia University Press, New York, 2013.
- [4] D. Laney. 3d data management: controlling data volume, velocity, and variety. Technical report, META Group, 2001. See gtnr.it/1bKfKH (PDF).
- [5] M. Li and P. Vitányi. *An Introduction to Kolmogorov Complexity and Its Applications*. Springer, New York, 2009.
- [6] National Research Council of the National Academies. *Frontiers in Massive Data Analysis*. The National Academies Press, Washington DC, 2013. ISBN-13: 978-0-309-28778-4. Online edition: bit.ly/14A0eyo.
- [7] J. G. Wolff. The management and analysis of big data, and the SP theory of intelligence. In preparation.
- [8] J. G. Wolff. Learning syntax and meanings through optimization and distributional analysis. In Y. Levy, I. M. Schlesinger, and M. D. S. Braine, editors, *Categories and Processes in Language Acquisition*, pages 179–215. Lawrence Erlbaum, Hillsdale, NJ, 1988. See bit.ly/ZIGjyc.

- [9] J. G. Wolff. Medical diagnosis as pattern recognition in a framework of information compression by multiple alignment, unification and search. *Decision Support Systems*, 42:608–625, 2006. See bit.ly/XE7pRG.
- [10] J. G. Wolff. *Unifying Computing and Cognition: the SP Theory and Its Applications*. CognitionResearch.org, Menai Bridge, 2006. ISBNs: 0-9550726-0-3 (ebook edition), 0-9550726-1-1 (print edition). Distributors, including Amazon.com, are detailed on bit.ly/WmB1rs.
- [11] J. G. Wolff. Application of the SP theory of intelligence to the understanding of natural vision and the development of computer vision. 2013. In preparation. See bit.ly/Xj3nDY (PDF).
- [12] J. G. Wolff. Information compression, intelligence, computing, and mathematics. 2013. In preparation. See bit.ly/1jEoECH.
- [13] J. G. Wolff. The sp theory of intelligence: an introduction. 2013. See bit.ly/1cFYTfw.
- [14] J. G. Wolff. The SP theory of intelligence: an overview. *Information*, 4(3):283–341, 2013. See bit.ly/19MmbLd.
- [15] J. G. Wolff. The SP theory of intelligence: benefits and applications. *Information*, 5(1):1–27, 2014. See bit.ly/1lcquWF.
- [16] G. K. Zipf. *Human Behaviour and the Principle of Least Effort*. Hafner, New York, 1949. Republished by Martino Publishing, Mansfield Centre, CT, 2012.